Shortcut Switching Enhancements for NHRP in DMVPN Networks

Last Updated: December 14, 2011

Routers in a Dynamic Multipoint VPN (DMVPN) network use Next Hop Resolution Protocol (NHRP) to discover the addresses of other routers and networks behind those routers that are connected to a nonbroadcast multiaccess (NBMA) DMVPN. The shortcut switching enhancements for NHRP provide an Address Resolution Protocol (ARP)-like solution that alleviates NBMA network problems, such as hub failure, decreased reliability, and complex configurations. With NHRP, systems attached to an NBMA network dynamically learn the NBMA address of the other systems that are part of that network, allowing these systems to directly communicate without requiring traffic to use an intermediate hop.

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Shortcut Switching Enhancements for NHRP

The following restrictions apply to this feature:

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

- The Shortcut Switching Enhancements for NHRP in DMVPN Networks feature is currently available only on the Cisco 870 series, Cisco 1600 series, Cisco 1700 series, Cisco 1800 series, Cisco 2600 series, Cisco 2800 series, Cisco 3600 series, Cisco 3700 series, Cisco 3800 series, Cisco 7200 series and Cisco 7301 routers. It specifically is not available for the Catalyst 6500 or Cisco 7600 series routers.
- Do not use this feature if DMVPN is configured with a partial full-mesh configuration; that is, if the router is configured with IP next-hop to be the IP address of the other spoke, then this feature is not required and must not be configured. In a full-mesh configuration the spokes are populated with a full routing table with next-hop being the other spokes.

Information About Shortcut Switching Enhancements for NHRP

NHRP in DMVPN Networks Overview

In previous implementations of DMVPN, the hub uses NHRP to maintain a database of the spokes' real (publicly reachable) IP addresses. Spokes in a DMVPN network register their real IP address with the hub using periodic NHRP registration packets. When a spoke has traffic for a destination behind another spoke, it uses an NHRP resolution request to query the NHRP database on the hub for the NBMA address of destination spokes. NHRP uses the resolution request process to build a direct spoke-to-spoke tunnel.

However, there were some issues with scaling implementations of DMVPN networks to large sizes (large number of spokes):

- In order for the spoke to "know" to send an NHRP resolution request to build a spoke-to-spoke tunnel, the spoke must have a route in its routing table for the remote network (behind the remote spoke) with an IP next-hop of the tunnel IP address of the remote spoke. Having each route in the routing table means that in a DMVPN network with 1000 spokes, each spoke router must have at least 1000 routes, one for each remote spoke. Having each route in the routing table also means that the hub router that helps in distributing this routing information must deal with the equivalent of 1,000,000 routes. The hub receives at least one route from each spoke, and the hub must advertise all of these routes to all of the other spokes; that is, 1000 routes advertised to 1000 spokes, which equals the processing of 1,000,000 routes. Supporting a routing table of that size can put a significant strain on the hub and routing protocol processing.
- When using the Open Shortest Path First (OSPF) routing protocol, OSPF must be configured to use broadcast network mode. Using broadcast network mode limits the number of hubs in a DMVPN network to just two. Only two hubs may be used because OSPF broadcast networks require a designated router (DR) and a backup designated router (BDR), and each spoke must be connected to both the DR and BDR. This configuration effectively limits the DMVPN network to two hubs when using OSPF.
- When using other routing protocols such as On-Demand Routing (ODR), which can only advertise a default-route 0.0.0.0/0 with an IP next-hop of the hub, the implementation could not support spoke-to-spoke tunnels.
- In order to scale DMVPN networks to larger sizes, multiple hubs are used where each one handles a subset of the spokes. For example, to handle 2000 spokes you could use four hubs, each one handling

500 spokes. The DMVPN implementation relied on hubs being “daisy-chained” together as NHRP Next Hop Servers (NHSs) of each other. In the case of daisy-chaining hubs, limitations included single hub failures breaking the chain, complexity in configuring multiple daisy chains to work around single hub failures, and delays in response time for building spoke-to-spoke tunnels as the chain of hubs grew.

- The DMVPN implementation only allowed a single hub spoke layer; hubs could not be spokes of other hubs within the same DMVPN network. You could not build complex DMVPN networks, such as having regional hubs that are spokes of central hubs. To build a similar design in previous implementations, you would have to set up different regional DMVPN networks with spokes connected to regional hubs and then interconnect the regional hubs outside of the DMVPN network or in a different DMVPN network. This design would mean that spokes in a region could only build spoke-to-spoke tunnels to spokes within the same region (same DMVPN network), but not to spokes in another region (different DMVPN network).
- In previous DMVPN implementations, the data packets were process switched at each hop along the daisy chain until the spoke-to-spoke tunnel was established. Process switching of data packets could result in unreasonable delays.

Benefits of NHRP Shortcut Switching Enhancements

The Shortcut Switching Enhancements for NHRP in DMVPN Networks feature provides a more scalable alternative to the previous NHRP model.

Cisco has developed NHRP shortcut switching model enhancements that allow for more scalable DMVPN implementations. This model provides the following advantages over previous DMVPN implementations:

- Allows summarization of routing protocol updates from hub to spokes. The spokes no longer need to have an individual route with an IP next-hop of the tunnel IP address of the remote spoke for the networks behind all the other spokes. The spoke can use summarized routes with an IP next-hop of the tunnel IP address of the hub and still be able to build spoke-to-spoke tunnels. It can reduce the load on the routing protocol running on the hub router. You can reduce the load because, when you can summarize the networks behind the spokes to a few summary routes or even one summary route, the hub routing protocol only has to advertise the few or one summary route to each spoke rather than all of the individual spoke routes. For example, with 1000 spokes and one router per spoke, the hub receives 1000 routes but only has to advertise one summary route to each spoke (equivalent to 1000 advertisements one per spoke) instead of the 1,000,000 advertisements it had to process in the prior implementation of DMVPN.
- Provides better alternatives to static daisy chaining of hubs for expanding DMVPN spoke-to-spoke networks. The hubs must still be interconnected, but they are not restricted to just a daisy-chain pattern. The routing table is used to forward data packets and NHRP control packets between the hubs. The routing table allows efficient forwarding of packets to the correct hub rather than having request and reply packets traversing through all of the hub routers.
- Allows for expansion of DMVPN spoke-to-spoke networks with OSPF as the routing protocol beyond two hubs. Because the spokes can use routes with the IP next-hop set to the hub router (not the remote spoke router as before), you can configure OSPF to use point-multipoint network mode rather than broadcast network mode. Configuring OSPF to use point-multipoint network mode removes the DR and BDR requirements that restricted the DMVPN network to just two hubs. When using OSPF, each spoke still has all individual routes, because the DMVPN network must be in a single OSPF area but you cannot summarize routes within an OSPF area.
- Allows routing protocols such as ODR to be used and still retain the ability to build dynamic spoke-to-spoke tunnels.
- Allows for hierarchical (greater than one level) and more complex tree-based DMVPN network topologies. Tree-based topologies allow capability to build DMVPN networks with regional hubs that

are spokes of central hubs. This architecture allows the regional hub to handle the data and NHRP control traffic for its regional spokes, but still allows spoke-to-spoke tunnels to be built between any spokes within in the DMVPN network, whether they are in the same region or not.

- Allows data packets to be Cisco Express Forwarding (CEF) switched along the routed path until a spoke-to-spoke tunnel is established.

NHRP Mapping and Adjacency Override

NHRP shortcut switching is now a feature in the CEF output feature switching path. For each data packet that is forwarded out the multipoint Generic Routing Encapsulation (mGRE) interface, NHRP performs a lookup in its mapping table to find an entry for the destination IP address of the data packet. If there is one, it overrides the adjacency determined by CEF during the Forwarding Information Base/Adjacency (FIB/ADJ) lookup. This lookup process is how data packets are redirected over the spoke-to-spoke direct tunnel rather than being forwarded to the hub as the routing table states.

If there is not a matching entry in the NHRP mapping table then the data packet is forwarded to the IP next-hop (adjacency) from the routing table; this would be the hub router. When this packet is received on the hub and it detects that this data packet has been received on and forwarded out the same tunnel interface, the hub router sends an NHRP redirect message to the previous tunnel hop (spoke router). When the spoke router receives the NHRP redirect, it sends an NHRP resolution request for the data packet destination IP address that triggered the NHRP redirect message. The NHRP resolution request and reply messages build a spoke-to-spoke tunnel between the two spokes behind which the hosts that are communicating are located. Once the spoke-to-spoke tunnel is built, an NHRP mapping entry is created to redirect the data packets over the spoke-to-spoke tunnel. If for some reason the spoke-to-spoke tunnel cannot be built, the data packets will continue to be forwarded via the hub(s).

For each NHRP mapping entry, NHRP keeps a reference to the CEF adjacency entry. This adjacency overrides the FIB-adjacency during CEF output feature processing. The figure below shows an example of the NHRP mapping table for shortcut switching.

Note To see if packets are being redirected over a spoke-to-spoke tunnel, you must look in the NHRP mapping table. The routing table and CEF FIB table will still show the original IP next-hop address.

Figure 1 NHRP Shortcut Switching Mapping Tables

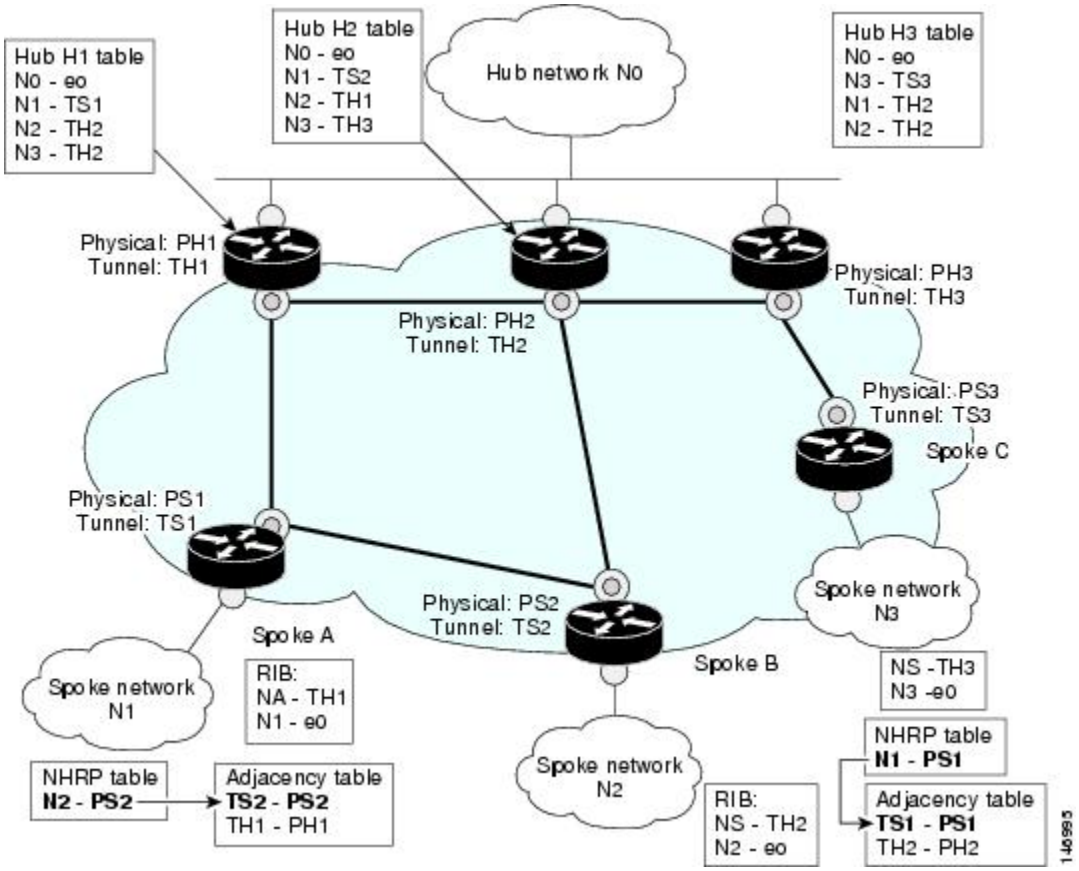

In the figure above, the packet flow is as follows:

1. The Spoke A forwarding table lookup for N2 gives TH1 as its next-hop and adjacency.
2. NHRP in the output feature path at Spoke A performs a lookup in its mapping table and does not find an entry for the host in N2.
3. The data packet is forwarded to Hub 1 (H1) as determined in Step 2.
4. H1 follows Steps 2 and 3 and forwards the data packet to Spoke B. NHRP in the output feature path also determines that the inbound (such as Tunnel0) and the outbound (Tunnel0) interface is part of the same DMVPN network and sends an NHRP redirect traffic indication to the tunnel (Spoke A) on which the data packet was received. The NHRP redirect message includes the original IP address and first eight bytes of the data packet.
5. At Spoke B the data packet is forwarded to the original host in network N2.
6. Spoke A receives an NHRP redirect traffic indication message from H1.
7. Spoke A processes the redirect message and triggers an NHRP resolution request for the destination IP address (host in N2) of the data packet (contained in the redirect message). NHRP will trigger a resolution only if it passes the NHRP interest list configuration check.

8 The NHRP resolution request follows the Spoke A-Hub 1-Spoke B path. The NHRP resolution follows the routed path towards the destination until it reaches Spoke B.
9 Spoke B routing lookup for the destination IP address (host in N2) finds that it is the exit point of the DMVPN network.
10 Spoke B builds any IPsec tunnel required to Spoke A and sends the NHRP resolution reply directly over the tunnel to Spoke A. The Spoke B reply contains prefix information of N2 and not just the reply for the host in N2; that is, Spoke B indicates that not only Host B but the entire network N2 is reachable via Spoke B. Spoke B creates a local cache entry for N2 and a list of requestors to which it has replied for network N2.
11 Spoke A receives the reply and installs the mapping for N2 in its mapping table. Further packets for any host in network N2 is forwarded to Spoke B directly.

NHRP Purge Request Reply

When an NHRP hub replies to a resolution request, it creates a local NHRP mapping entry. The local mapping entry is a network entry for which NHRP has sent a reply. The local mapping entry maintains a list of requestors. When a network entry is modified or deleted in the routing table, NHRP is notified of the event. NHRP finds the local cache entry for the network and sends a purge request to the requestors that the network to which it previously replied has changed. The receivers of the purge message delete the corresponding NHRP mapping entry from its table and send a purge reply indicating that the purge message was processed successfully.

How to Configure Shortcut Switching for NHRP

Note If **ip nhrp shortcut** and **ip nhrp redirect** are not configured, then the DMVPN network will continue to function as it did prior to this feature.

- Enabling NHRP Shortcut Switching on an Interface, page 6
- Configuring NHRP Redirect, page 7

Enabling NHRP Shortcut Switching on an Interface

Perform this task to enable shortcut switching for NHRP for an interface on a router.

Note When using this feature, we recommend configuring the **ip nhrp redirect** command on all the DMVPN nodes. This configuration would be useful in the event the data traffic takes a spoke-to-spoke-hub-spoke path.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **ip nhrp shortcut**
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**interface** *type number* **Example:** `Router(config)# interface Tunnel0`	Enters interface configuration mode.
Step 4	**ip nhrp shortcut** **Example:** `Router(config-if)# ip nhrp shortcut`	Enables NHRP shortcut switching on an interface.
Step 5	**end** **Example:** `Router(config-if)# end`	Ends the configuration session.

Configuring NHRP Redirect

NHRP sends a resolution request for a shortcut path after receiving an NHRP redirect traffic indication message. An NHRP redirect traffic indication is generated by an intermediate node when a data packet is forwarded within the same DMVPN network (in and out the same tunnel interface). The redirect is sent to the previous tunnel hop (spoke) on the tunnel from which the data packet was received.

The NHRP redirect traffic indication is generated for each unique combination of source-NBMA IP address (previous tunnel hop) and data packet (destination IP address); that is, redirect is generated independent of the source IP address of the data packet. It totally depends on the destination IP address and the source-NBMA address of the incoming Generic Routing Encapsulation (GRE) encapsulated data packet. These NHRP redirect messages are rate-limited. A configurable option is provided to determine the rate at which NHRP redirects will be generated for the same combination of source-NBMA address and data destination IP address.

Like an Internet Control Message Protocol (ICMP) message, the NHRP redirect message includes the IP header and the first eight data bytes of the data packet that triggers the redirect. This information is used by NHRP on the previous tunnel hop to determine whether and where to send a resolution request. That is, NHRP would match against the interest list configuration to determine whether to send a resolution request.

Perform this task to enable NHRP redirects.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface tunnel** *number*
4. **ip nhrp redirect [** every **{seconds}]**
5. end
6. **show ip nhrp traffic**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**interface tunnel** *number* **Example:** `Router(config)# interface tunnel0`	Enters tunnel interface configuration mode.

	Command or Action	Purpose
Step 4	**ip nhrp redirect [** every **{seconds}]** **Example:** `Router(config-if)# ip nhrp redirect every 1`	Enables redirect traffic indication if traffic is forwarded with the NHRP network. Use the **every** keyword and *seconds* argument to indicate when to expire a redirect entry created to avoid sending duplicate redirects.
Step 5	end **Example:** `Router(config-if)# end`	Ends the configuration session.
Step 6	**show ip nhrp traffic** **Example:** `Router# show ip nhrp traffic`	(Optional) Displays NHRP traffic statistics. • This command will display the number of NHRP traffic indication packets (redirects) originated from or received by the station.

Configuration Examples for Shortcut Switching Enhancements for NHRP

Configuring NHRP Shortcut Switching and NHRP Redirect Example

The following example shows how to configure NHRP shortcut switching and NHRP redirect on tunnel interface 0:

```
Router> enable

Router# configure terminal
Router(config)# interface Tunnel0
Router(config-if)# ip address 192.2.0.11 255.255.255.0
Router(config-if)# ip nhrp authentication test
Router(config-if)# ip nhrp map multicast 192.2.0.2
Router(config-if)# ip nhrp map 192.2.0.2 192.2.0.13
Router(config-if)# ip nhrp network-id 100000
Router(config-if)# ip nhrp nhs 192.2.0.11
Router(config-if)# ip nhrp shortcut
Router(config-if)# ip nhrp redirect
Router(config-if)# tunnel source Serial1/0
Router(config-if)# tunnel mode gre multipoint
Router(config-if)# tunnel key 100000
Router(config-if)# tunnel protection ipsec profile vpnprof
```

Additional References

Related Documents

Related Topic	Document Title
NHRP information and configuration tasks	"Configuring NHRP" module
NHRP commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	*Cisco IOS IP Addressing Services Command Reference*
Dynamic Multipoint VPN	"Dynamic Multipoint VPN" module

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	--

MIBs

MIB	MIBs Link
None	--

RFCs

RFC	Title
None	--

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Shortcut Switching Enhancements for NHRP in DMVPN Networks

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1 Feature Information for Shortcut Switching Enhancements for NHRP in DMVPN Networks

Feature Name	Releases	Feature Information
Shortcut Switching Enhancements for NHRP in DMVPN Networks	12.4(6)T	Routers in a Dynamic Multipoint VPN (DMVPN) network can use the Next Hop Resolution Protocol (NHRP) to discover the addresses of other routers and networks behind those routers that are connected to a DMVPN nonbroadcast multiaccess (NBMA) network. NHRP provides an ARP-like solution that alleviates NBMA network problems, such as hub failure, decreased reliability, and complex configurations. The following commands were introduced or modified: **clear ip nhrp shortcut, debug dmvpn, debug nhrp routing, ip nhrp shortcut**, **show dmvpn**, **show ip nhrp, show ip nhrp shortcut**, **show ip route**, **show ip route next-hop-override**.